

Information-theoretic limits of Bayesian network structure learning

Asish Ghoshal and Jean Honorio
 Department of Computer Science
 Purdue University
 West Lafayette, IN - 47906
 {aghoshal, jhonorio}@purdue.edu

In this paper, we study the information-theoretic limits of learning the structure of Bayesian networks (BNs), on discrete as well as continuous random variables, from a finite amount of data. We show that the minimum number of samples required by *any procedure* to recover the correct structure grows as $\Omega(m)$ and $\Omega(k \log m + k^2/m)$ for non-sparse and sparse BNs respectively, where m is the number of variables and k is the maximum number of parents per node. We provide a simple recipe, based on an extension of the Fano’s inequality, to obtain information-theoretic limits of structure recovery for any exponential family BN. We instantiate our result for specific conditional distributions in the exponential family to characterize the fundamental limits of learning various commonly used BNs, such as Conditional Probability Table (CPT) based networks, Gaussian BNs, Noisy-OR networks, and Logistic regression (LR) networks. For each of the aforementioned networks, we identify important parameters of the conditional distributions that affect the complexity of learning such models. We also show that an existing procedure called *SparsityBoost*, for learning binary CPT networks is not information-theoretically optimal in the number of variables. En route to obtaining our main results, we obtain tight bounds on the number of sparse and non-sparse essential-DAGs. Finally, as a byproduct, we recover the information-theoretic limits of sparse variable selection for logistic regression.

1 Introduction

Motivation. Bayesian Networks (BNs) are a class of probabilistic graphical models that describe the conditional dependencies between a set of random variables as a directed acyclic graph (DAG). However, in many problems of practical interest, the structure of the network is not known a priori and must be inferred from data.

The problem of BN structure learning has received increased attention in the past several years. Several algorithms have been developed for consistently recovering the best BN, in terms of “closeness” to the true distribution or true structure, for some measure of closeness, from a given sample of data; for a detailed survey of various algorithms see, e.g., [1] and [2]. However, an important question that has hitherto remained unanswered is the fundamental limits of learning

BNs, i.e., “What is the minimum number of samples required by any procedure to recover the true DAG structure of a BN?”. The answer to this question would help shed light on the fundamental limits of learning the DAG structure of BNs, and also help determine if existing algorithms are optimal in terms of their sample complexity or if there exists a gap between the state-of-the-art estimation procedures and the information-theoretic limits. In this paper we obtain lower bounds on the minimum number of samples required to learn BNs over m variables, and sparse BNs over m variables with maximum in-degree of k .

Comparison with Markov Random Fields. While there has been a lot of prior work on determining the information-theoretic limits of structure recovery in Markov Random Fields (MRFs), which are undirected graphical models, characterizing the information-theoretic limits of learning BNs (directed models) is important in its own right for the following reasons. First, unlike MRFs where the undirected graph corresponding to a dependence structure is uniquely determined, multiple DAG structures can encode the same dependence structure in BNs. Therefore, one has to reason about Markov equivalent DAG structures in order to characterize the information-theoretic limits of structure recovery in BNs. Second, the complexity of learning MRFs is characterized in terms of parameters of the joint distribution over nodes, which in turn relates to the overall graph structure, while the complexity of learning BNs is characterized by parameters of local conditional distributions of the nodes. The latter presents a technical challenge, as we show later, when the marginal or joint distribution of the nodes in a BN do not have a closed form solution.

Contribution. The main contributions of this paper are as follows:

- (I) We derive necessary conditions on the sample complexity of recovering the DAG structure of non-sparse and sparse BNs. We show that $\Omega(m)$ samples are necessary for consistent recovery of the DAG structure of BNs, while for sparse networks $\Omega(k \log m + k^2/m)$ samples are necessary.
- (II) We provide a simple recipe for obtaining the information-theoretic limits of learning any exponential family BN. We instantiate our result for specific conditional distributions to determine the fundamental limits of learning the structure of various widely used BNs, namely, conditional probability table (CPT) based networks, Gaussian networks, noisy-OR networks, and logistic regression networks.
- (III) We show that the *SparsityBoost* algorithm developed by Brenner and Sontag [3] for learning binary CPT BNs, which has a sample complexity of $\mathcal{O}(m^2(1/\theta_{\min}))$, is far from the information-theoretic limit of $\Omega\left(\frac{k \log m + k^2/m}{\log(1/\theta_{\min})}\right)$, where θ_{\min} , is the minimum probability value in the conditional probability tables.
- (IV) We are able to demonstrate optimality of ℓ_1 -regularized logistic regression for support recovery. The necessary conditions of $\Omega(k^2 \log m)$ derived in this paper, matches the sufficient condition of $\mathcal{O}(k^2 \log m)$, obtained by Ravikumar et al. [4].
- (V) Finally, an interesting corollary of our main result is that learning *layered BNs* — where the ordering of the nodes is known (upto layers) and the parent set of each variable is constrained to be in the layer above it — is as hard as learning general BNs in terms of their sample complexity.

Of lesser interest, yet interesting, is our comparison of the complexity of learning Gaussian BNs with Gaussian MRFs. We obtain slightly stronger lower bounds for learning Gaussian BNs than the corresponding bounds for Gaussian MRFs. Lastly, of independent interest, are our extension of the Fano’s inequality to the case where there are latent variables, and our upper-bound on the KL divergence between two exponential family distributions as the inner product of the difference between the natural parameter and expected sufficient statistics.

The rest of the paper is organized as follows. In Section 2, we discuss relevant prior work on structure learning of BNs. In Sections 3 and 4, we present some background information about BNs and formalize the problem of structure recovery in BNs, respectively. Then, we present our main results for learning BNs in Section 5, which we then instantiate for specific conditional distributions to obtain sample complexity results for various widely used networks in Section 6. Finally, we conclude with some directions for future work in Section 7.

2 Related Work

Höffgen [10], and Friedman and Yakhini [11] were among the first to derive sample complexity results for learning BNs. In both [10] and [11] the authors provide upper bounds on sample complexity of learning a BN that is likelihood consistent, i.e. the likelihood of the learned network is ϵ away from the likelihood of the true network in terms of the Kullback-Leibler (KL) divergence measure. In other words, [10] and [11] do not provide guarantees for recovering the true structure of the network. Moreover, in [10] the authors consider boolean variables and each variable is allowed to have at most k parents. Abbeel et al. in [12] provide polynomial sample complexity results for learning likelihood consistent factor graphs.

Among sample complexity results for learning structure consistent BNs, where the structure of the learned network is close to the true network, Spirtes et al. [13] and Cheng et al. [14] provide such guarantees for polynomial-time test-based methods, but the results hold only in the infinite-sample limit. Chickering and Meek in [15] also provide a greedy hill-climbing algorithm for structure learning that is structure consistent in the infinite sample limit. Zuk et al. [16] show structure consistency of a single network and do not provide uniform consistency for all candidate networks, i.e. the bounds relate to the error of learning a specific wrong network having a score greater than the true network. Brenner and Sontag [3] provide upper bounds on the sample complexity of recovering the structure of sparse BNs. However, they consider binary valued variables only and the sample complexity grows as $\mathcal{O}(m^2)$.

While there has been some prior work on obtaining the information-theoretic limits of learning graphical models, almost all of them focused on a specific class of undirected graphical models called Markov Random Fields (MRFs); whereas in this paper we focus on directed graphical models, i.e., BNs. A recurring theme in all such works is to construct ensembles of MRFs that are hard to learn and then use the Fano’s inequality to lower bound the estimation error by treating the inference procedure as a communication channel. Santhanam and Wainwright [6] obtained necessary and sufficient conditions for learning pairwise binary MRFs. The necessary and sufficient conditions on the number of samples scaled as $\mathcal{O}(k^2 \log m)$ and $\mathcal{O}(k^3 \log m)$ respectively, where k is the maximum node degree. Information theoretic limits of learning Gaussian MRFs was studied by Wang et al. [7] and for walk-summable Gaussian networks, by Anandkumar et al. [17]. In [18], Anandkumar et al. obtain a necessary condition of $\Omega(c \log m)$ for structure learning of Erdős-Rényi random Ising models, where c is the average node degree.

3 Preliminaries

Let $\mathbf{X} = \{X_1, \dots, X_m\}$ be a set of random variables, where $X_i \in \mathcal{X}_i, \forall i \in [m]$. Let $\mathcal{D} \stackrel{\text{def}}{=} \times_{i=1}^m \mathcal{X}_i$ be the domain in which the variables in \mathbf{X} jointly take their values. A BN for \mathbf{X} is a tuple $(G, \mathcal{P}(G, \Theta))$; where $G = (V, E)$ is a directed acyclic graph (DAG) with $V = [m]$ being the vertex set and $E \subset [m] \times [m]$ being the set of directed edges, and $\mathcal{P}(G, \Theta)$ is a probability distribution over \mathbf{X} that is parameterized by Θ and factorizes according to the DAG structure G . Particularly, $\forall \mathbf{x} \in \mathcal{D}, \mathcal{P}(\mathbf{x}; G, \Theta) = \prod_{i=1}^m \mathcal{P}_i(x_i; \pi_i(G), \Theta)$, where $\pi_i(G) \subseteq [m] \setminus \{i\}$ is the parent set of the i -th node in G , $\mathbf{X}_{\pi_i(G)} = \{X_j | j \in \pi_i(G)\}$, $\mathcal{P}_i(\cdot) = \Pr\{x_i | \mathbf{X}_{\pi_i(G)}, \Theta_i\}$ is the conditional distribution of X_i given an assignment to its parent set, and Θ_i are the parameters for the i -th conditional distribution.

The DAG structure G of a BN specifies the conditional independence relationships that exist between different random variables in the set \mathbf{X} . Different graph structures which make the same conditional independence assertions about a set of random variables are called Markov equivalent.

Definition 1 (Markov equivalence). *Two DAGs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ are Markov equivalent if for all disjoint subsets $A, B, C \subset V$, $X_A \perp\!\!\!\perp X_B | X_C$ in $G_1 \iff X_A \perp\!\!\!\perp X_B | X_C$ in G_2 .*

The set of DAGs that are Markov equivalent to the DAG G is denoted by $[G]$.

An essential graph¹, consisting of both directed edges, which are called protected edges, and undirected edges, is a canonical representation of the (Markov) equivalence class of DAGs. The undirected edges can be oriented in either direction without changing the conditional independence relationships encoded by the graphs. We denote by G^* the essential graph for $[G]$.

4 Problem Formulation

Let \mathcal{G} be an ensemble of DAGs. We denote by \mathcal{G}_m the ensemble of DAGs over m nodes. Also, let $\varphi(\mathcal{G})$ be some set of “parameter maps”. A parameter map $\Theta \in \varphi(\mathcal{G})$, maps a given DAG structure G to a specific instance of the conditional distribution parameters that are compatible with the DAG structure G , i.e. $\Theta(G)$. It is useful to think of Θ as a policy for setting the parameters of the conditional distributions, given a DAG G . For instance, for binary CPT networks, a particular policy Θ would consist of several candidate probability tables for each node, one for each possible number of parents the node can have (from 0 to $m - 1$), with entries set to some specific values. Then, given a DAG structure G , $\Theta(G)$ assigns a probability table to each node (from the policy Θ) according to the number of parents of the node in G . This notion of parameter maps affords us the ability to generate a BN by sampling the DAG structure and the parameters independently of each other, which, as would be evident later, is a key technical simplification. Let $\mathcal{P}_{\mathcal{G}}$ and $\mathcal{P}_{\varphi(\mathcal{G})}$ be probability measures on the set \mathcal{G} and $\varphi(\mathcal{G})$ respectively. Nature picks a graph structure G , according to $\mathcal{P}_{\mathcal{G}}$, and then samples a parameter map Θ , independently, according to $\mathcal{P}_{\varphi(\mathcal{G})}$. Thereafter, nature generates a data set of n i.i.d observations, $\mathbf{S} = \{\mathbf{x}^{(i)}\}_{i=1}^n$, with $\mathbf{x}^{(i)} \in \mathcal{D}$, from the BN $(G, \mathcal{P}(G, \Theta(G)))$. The problem of structure learning in BNs concerns with estimating the graph structure \hat{G} , up to Markov equivalence, from the data set \mathbf{S} . In that context, we define the notion of a decoder. A decoder is any function $\zeta : \mathcal{D}^n \rightarrow \mathcal{G}$ that maps a data set of n observations to an estimated DAG $\hat{G} \in \mathcal{G}$. The estimation error is defined as follows

$$p_{\text{err}} \stackrel{\text{def}}{=} \inf_{\zeta} \max_{\mathcal{P}_{\varphi(\mathcal{G})}} \mathcal{P}_{\varphi(\mathcal{G})}(\Theta) \mathcal{P}_{\mathcal{G}}(G) \Pr\{\zeta(\mathbf{S}) \notin [G] | G, \Theta\}, \quad (1)$$

¹See Andersson et. al. [5] for a formal definition.

where the probability $\Pr \{ \cdot | G, \Theta \}$ is computed over the data distribution $\mathcal{P}(G, \mathcal{P}(G, \Theta(G)))$ for a specific DAG structure G and parameters $\Theta(G)$.

Note that our definition of estimation error is stronger than what is typically used in the literature for structure recovery of MRFs (see e.g. [6] and [7]), since we focus on the maximum error across all measures $\mathcal{P}_{\varphi(G)}$ over the parameter maps $\varphi(G)$ which itself can be uncountable. Here, we are interested in obtaining necessary conditions for consistent structure recovery of BNs, i.e., we show that if the number of samples is less than a certain threshold, then any decoder ζ fails to recover the true graph structure with probability of error $p_{\text{err}} > 1/2$.

We emphasize that while our sample complexity results invariably depend on the parameter space $\varphi(\mathcal{G})$ under consideration, the decoder only has access to the data set \mathbf{S} . Apart from \mathcal{G}_m , we consider various other ensembles of DAGs in this paper, to fully characterize the fundamental complexity of learning different classes of BNs. Among the ensembles we consider, $\mathcal{G}_{m,k}$ denotes the family of DAGs, where each node is allowed to have at most k parents. We also consider generalizations of QMR-DT [8] type two-layered BNs, to multiple layers of nodes, with nodes in each layer only allowed to have parents in the layer above it.

Let $\mathcal{V} = \{\mathbf{V}_i\}_{i=1}^l$, define an ordering of m nodes into l layers where \mathbf{V}_i is the set of nodes in the i -th layer. We have that $|\mathbf{V}_i| = m_i$ and $\sum_{i=1}^l |\mathbf{V}_i| = m$. $\mathcal{G}_m^l(\mathcal{V})$ denotes an ensemble of DAG structures where $\forall G = (\mathbf{V}, \mathbf{E}) \in \mathcal{G}_m^l$, $\mathbf{V} = \bigcup_{i \in \mathcal{V}} \mathbf{V}_i$ and $\mathbf{E} = \{(u, v) | u \in \mathbf{V}_{i+1} \wedge v \in \mathbf{V}_i, i \in [l-1]\}$. We write \mathcal{G}_m^l instead of $\mathcal{G}_m^l(\mathcal{V})$ to indicate that the members of \mathcal{G}_m^l have some *known* layer-wise ordering of the m nodes, without making the ordering explicit. Finally, we consider another ensemble $\mathcal{G}_{m,k}^l \subset \mathcal{G}_m^l$ where the nodes are allowed to have at most k parents. Together, the ensembles $\mathcal{G}_m, \mathcal{G}_{m,k}, \mathcal{G}_m^l$ and $\mathcal{G}_{m,k}^l$, span a wide range of the sample complexity landscape of recovering the structure of BNs. In the following section we present our main result on the fundamental limits of learning BNs.

5 Main Results

Fano’s inequality is one of the primary tools used for deriving necessary conditions on structure recovery of graphical models. The difficulty of recovering the DAG structure of a BN, however, depends both on the structural properties of the ensemble of DAG structures under consideration, as well as on the conditional distributions and their parameters. In order to obtain guarantees about structure recovery, we treat the parameters of the conditional distributions as latent variables — variables that we do not observe and are not interested in estimating. Given that the likelihood of the observed data depends, both on the structure and parameters of the BN that generated the data, it behooves us to ask: “*If we are only interested in recovering the structure of BNs, do the presence of unobserved parameters make structure estimation easier or harder?*” To rigorously answer this question, we extend the classic Fano’s inequality, which is defined for a Markov chain, to a slightly more general setting as given below.



Figure 1: Fano’s inequality extension. In (a) the edge between W and X is undirected to indicate that the edge can be oriented in either direction.




Figure 2: The DAG structure G and a parameter map Θ are sampled independently. Data set S of m samples is generated from the BN $(G, \mathcal{P}(G, \Theta(G)))$. A decoder ζ then estimates the DAG structure $\hat{G} = \zeta(S)$.

Theorem 1 (Fano’s inequality extension). *Let W, X , and Y be random variables and let \hat{X} be any estimator of X . If the random variables are related according to the graphical model in Figure 1 (a), then*

$$\Pr \{X \neq \hat{X}\} \geq 1 - \frac{I(Y; X|W) + \log 2}{H(X|W)}. \quad (2)$$

Moreover, if $W \in \mathcal{W}$, is independent of X (Figure 1 (b)), and $\mathcal{P}_{\mathcal{W}}$ be any probability measure over \mathcal{W} then,

$$\max_{\mathcal{P}_{\mathcal{W}}} \mathcal{P}_{\mathcal{W}}(w) \Pr \{X \neq \hat{X}|W = w\} \geq 1 - \frac{\sup_{w \in \mathcal{W}} I(Y; X|W = w) + \log 2}{H(X)}. \quad (3)$$

Missing proofs are included in the Appendix.

Remark 1. *Theorem 1 can be seen an extension of Fano’s inequality to the case where there are latent variables W that influence Y , while we are interested in only estimating X . If in Figure 1 we have that $W \rightarrow X$, then $I(X; Y|W) \geq I(X; Y)$. Further, since $H(X|W) \leq H(X)$, we conclude that presence of the latent variable W , reduces the estimation error p_{err} .*

Remark 2. *When X and W are independent, we get $I(X; Y|W) = I(X; Y)$, and (2) reduces to the well known Fano’s inequality. However, the conditional mutual information $I(X; Y|W = w)$ in 3 can be computed easily as compared to $I(X; Y)$ when we have access to the conditional distribution of Y given X and W . Also note that we do not need \mathcal{W} to be countable.*

Theorem 1 serves as our main tool for lower bounding the estimation error, and subsequently obtaining necessary conditions on the number of samples. In order to obtain sharp lower bounds on p_{err} , we assume that DAG structures and parameters maps (and by extension parameters) are sampled independently. Figure 2 shows the schematics of the inference procedure.

Henceforth, we will use the terms “parameter maps” and “parameters” interchangeably, since given a DAG structure G , the parameter map Θ , maps G to a specific parameterization of the BN. Then, given any ensemble of DAG structures \mathcal{G} , the main steps involved in using Theorem 1 to lower bound p_{err} are: (a) obtaining lower bounds on $H(G) = \log |\mathcal{G}|$, which follows from our assumption that G is sampled uniformly from \mathcal{G} , and (b) computing the mutual information between the data set and DAG structures over all possible parameter choices, i.e., $\sup_{\Theta \in \varphi(\mathcal{G})} I(S; G|\Theta)$. To accomplish each of the above objectives, we consider restricted ensembles $\tilde{\mathcal{G}}_m \subset \mathcal{G}_m$ and $\tilde{\mathcal{G}}_{m,k} \subset \mathcal{G}_{m,k}$, of size-one equivalence classes, i.e., $\forall G \in \tilde{\mathcal{G}}_m \vee G \in \tilde{\mathcal{G}}_{m,k}, |[G]| = 1$. Note that for any graph G in \mathcal{G}_m^l or $\mathcal{G}_{m,k}^l$, we have that $|[G]| = 1$, since edges are constrained to go from layer $(i + 1)$ to i . Thus, the ensembles $\tilde{\mathcal{G}}_m, \tilde{\mathcal{G}}_{m,k}, \mathcal{G}_m^l$ and $\mathcal{G}_{m,k}^l$ can be thought of as consisting of essential DAGs, where all edges are protected. In the following section, we bound the number of essential DAGs in each of the restricted ensembles.

Enumerating DAGs. Essential DAGs, i.e., Markov equivalent classes of DAGs of size 1, was first enumerated by Steinsky [9]. However, the number of DAGs is given as a recurrence relation, for which a closed form solution is difficult to compute. Therefore, we compute tight bounds on the number of essential DAGs in the following paragraphs. In the following lemmas we bound the number of DAGs in each of the restricted ensembles introduced previously.

Lemma 1. *The size of the restricted ensemble $\tilde{\mathcal{G}}_m$ is bounded as follows:*

$$2^{(m(m-3)/2)+1} \leq |\tilde{\mathcal{G}}_m| \leq m! 2^{m(m-1)/2}, \quad (4)$$

and $\log |\tilde{\mathcal{G}}_m| \geq ((m(m-3)/2) + 1) \log 2$.

Note that the lower bound in Lemma 1 is asymptotically tight. Now we bound the number of essential DAGs where each node is allowed to have at most k parents. The proof the following lemma is also based on the inclusion-exclusion method used by Steinsky, and can be found in the Appendix.

Lemma 2. *Assuming $k > 1$ and $m > 2$, the size of the restricted ensemble $\tilde{\mathcal{G}}_{m,k}$ is bounded as follows:*

$$2^{(k(k-3)/2)+1} \prod_{j=k+1}^{m-1} \left(\sum_{i=0}^k \binom{j-1}{i} \right) \leq |\tilde{\mathcal{G}}_{m,k}| \leq m! 2^{k(k-1)/2} \prod_{j=k+1}^{m-1} \left(\sum_{i=0}^k \binom{j}{i} \right), \quad (5)$$

and $\log |\tilde{\mathcal{G}}_{m,k}| \geq k \{ \log(m-2)! - \log k! - (m-k-2) \log k \} + \{ (k(k-3)/2) + 1 \} \log 2$.

Note that using Stirling's factorial formula, the above lemma gives the following lower bound on the number of sparse essential DAGs: $\log |\tilde{\mathcal{G}}_{m,k}| = \Omega(km \log m)$. A little calculation shows that $\log |\tilde{\mathcal{G}}_{m,k}| = \mathcal{O}(km \log m)$ for large enough m . Thus our bounds for the number of sparse essential DAGs is tight. The following lemma bounds the number of “layered” essential DAGs.

Lemma 3. *The number of BNs in the family \mathcal{G}_m^l and $\mathcal{G}_{m,k}^l$ is as follows:*

$$|\mathcal{G}_m^l| = \prod_{i=1}^{l-1} (2^{m_{i+1}})^{m_i}, \quad |\mathcal{G}_{m,k}^l| = \prod_{i=1}^{l-1} \left[\sum_{j=0}^k \binom{m_{i+1}}{j} \right]^{m_i}. \quad (6)$$

Further, $\log |\mathcal{G}_m^l|$ and $\log |\mathcal{G}_{m,k}^l|$ are given as follows:

$$\log |\mathcal{G}_m^l| = (\log 2) \sum_{i=1}^{l-1} (m_{i+1})(m_i), \quad \log |\mathcal{G}_{m,k}^l| \geq k \sum_{i=1}^{l-1} m_i \log \left(\frac{m_{i+1}}{k} \right).$$

Next, we compute bounds on the mutual information between the data set and DAG structures.

Mutual Information Bounds. Computing the mutual information $I(\mathbf{S}; G | \Theta)$ exactly, is in general difficult. Therefore, we use the following lemma to bound the mutual information from above.

Lemma 4. Let $\mathcal{P}_{\mathbf{S}|G,\Theta}$ be the distribution of \mathbf{S} conditioned on a specific DAG G and specific parameters Θ , and let \mathcal{Q} be any distribution over \mathbf{S} . Then we have

$$\sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) \leq \sup_{\Theta \in \varphi(\mathcal{G})} \frac{1}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \mathbb{KL}(\mathcal{P}_{\mathbf{S}|G,\Theta} \| \mathcal{Q}). \quad (7)$$

Assuming that $\mathcal{X}_i = \mathcal{X}$, $\forall i \in [m]$, we chose the distribution \mathcal{Q} to be the product distribution $\mathcal{Q} = \mathcal{Q}_0^{mn}$, where $\mathcal{Q}_0 = \mathcal{P}_i(\emptyset, \Theta_i)$. In other words, the distribution \mathcal{Q} is chosen to be the distribution encoded by a DAG with no edges.

The main hurdle in using (7) to bound the mutual information $I(\mathbf{S}, G|\Theta)$, is computing the KL divergence $\mathbb{KL}(\mathcal{P}_{\mathbf{S}|G,\Theta} \| \mathcal{Q})$. Often times, in BNs characterized by local conditional distributions, coming up with a closed form solution for the joint distribution over all nodes or even the marginal distribution of an arbitrary node is not possible; unless we assume that the marginal distribution of the parents of a node form a *conjugate prior* for the conditional distribution of the node — an assumption which is quite restrictive. Therefore, to tackle the above problem, we derive the following upper bound on the KL divergence for exponential family distributions which is easy to compute.

Lemma 5 (KL Divergence Bound for Exponential Family Distributions). Let $X \in \mathbb{R}^d$ be any random variable. Let \mathcal{P}_1 and \mathcal{P}_2 be distributions over X , belonging to the exponential family, with natural parameters $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$ respectively, i.e.

$$\mathcal{P}_1(x) = \exp(\boldsymbol{\eta}_1^T \mathbf{T}(x) - \psi(\boldsymbol{\eta}_1)) h(x),$$

where $\mathbf{T}(X)$ is the sufficient statistics (similarly for \mathcal{P}_2). Assuming $h(x) \neq 0 \forall x \in \mathbb{R}^d$, we have

$$\mathbb{KL}(\mathcal{P}_1 \| \mathcal{P}_2) \leq \Delta(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2), \quad (8)$$

$$\Delta(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2) \stackrel{\text{def}}{=} (\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2)^T (\mathcal{T}(\boldsymbol{\eta}_1) - \mathcal{T}(\boldsymbol{\eta}_2)), \quad (9)$$

where $\mathcal{T}(\boldsymbol{\eta}_1) \stackrel{\text{def}}{=} \mathbb{E}_X[\mathbf{T}(x)|\boldsymbol{\eta}_1]$ is the expected sufficient statistic of X as computed by the distribution parameterized by $\boldsymbol{\eta}_1$ (similarly for $\mathcal{T}(\boldsymbol{\eta}_2)$).

Note that even though $\mathbb{KL}(\mathcal{P}_1 \| \mathcal{P}_2)$ is not symmetric, its upper bound $\Delta(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2)$ is symmetric. Given the fact that \mathbf{S} is sampled i.i.d from $\mathcal{P}(G, \Theta(G))$, which in turn factorizes as a product of conditional distributions \mathcal{P}_i , we then have the following result for the mutual information $I(\mathbf{S}; G|\Theta)$.

Lemma 6 (Mutual Information Bound). For any ensemble of DAG structures \mathcal{G} , we have

$$\sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) \leq \frac{n}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \sum_{i=1}^m \sup_{\Theta \in \varphi(\mathcal{G})} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\mathbb{KL}(\mathcal{P}_i(\pi_i(G), \Theta) \| \mathcal{Q}_0)], \quad (10)$$

where $\mathcal{P}_i(\pi_i(G), \Theta)$ is the conditional distribution of the i -th node. Further, if we have that, $\forall i \in [m], X_i \in \mathcal{X}$ and $\mathcal{P}_i(\pi_i(G), \Theta_i)$ belongs to the exponential family with natural parameter $\boldsymbol{\eta}_i \stackrel{\text{def}}{=} \boldsymbol{\eta}(\mathbf{X}_{\pi_i}, \Theta_i)$ and \mathcal{Q}_0 belongs to the exponential family with natural parameter $\boldsymbol{\eta}_0$; then,

$$\sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) \leq \frac{n}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \sum_{i=1}^m \sup_{\Theta \in \varphi(\mathcal{G})} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)]. \quad (11)$$

Remark 3. In the above lemma, $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$ is a random variable because the natural parameter $\boldsymbol{\eta}_i$ depends on the parents \mathbf{X}_{π_i} . The quantity $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$ in the above lemma is non-negative and measures how far the conditional distribution of a variable with parents $\pi_i(G)$ is from the distribution of the variable with no parents, as a function of the difference between the expected sufficient statistics and the natural parameters. The mutual information between the data set S and the DAG structure G is then a sum of the expected “distances” of the conditional distributions from the distribution of a variable with no parents.

With the exception of Gaussian BNs, where we can write the joint and marginal distributions of the variables in closed form, it is in general difficult to compute the expectation of $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$. Therefore, we bound the mutual information by bounding $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$, which can be easily done for bounded random variables. From the above lemma, we then get the following mutual information bound for layered BNs.

Corollary 1 (Mutual Information Bound for Layered BNs). *If $\mathcal{G} = \mathcal{G}_m^l(\mathcal{V}) \vee \mathcal{G} = \mathcal{G}_{m,k}^l(\mathcal{V})$, then*

$$\sup_{\Theta \in \boldsymbol{\varphi}(\mathcal{G})} I(S; G | \Theta) \leq \frac{(m - m_l)n}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \left\{ \max_{i \in \mathbf{V} \setminus \mathbf{V}_l} \sup_{\Theta \in \boldsymbol{\varphi}(\mathcal{G})} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\text{KL}(\mathcal{P}_i(\pi_i(G), \Theta) \| \mathcal{Q}_0)] \right\}, \quad (12)$$

where we recall that $\mathcal{V} = \{\mathbf{V}_j\}_{j=1}^l$ is an ordering of nodes into l layers, \mathbf{V}_j is the set of nodes in the j -th layer and $m_j = |\mathbf{V}_j|$. Further, for exponential family conditional distributions, we have

$$\sup_{\Theta \in \boldsymbol{\varphi}(\mathcal{G})} I(S; G | \Theta) \leq \frac{(m - m_l)n}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \left\{ \max_{i \in \mathbf{V} \setminus \mathbf{V}_l} \sup_{\Theta \in \boldsymbol{\varphi}(\mathcal{G})} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)] \right\}. \quad (13)$$

In order to obtain tight sample complexity results, we need to create “difficult instances” of BNs that are hard to learn. Intuitively speaking, inferring the parents of a node will be hard if the conditional distribution of a node with many parents is close to that of a node with no parents. Therefore, we make the following crucial assumption about the conditional distributions specified by the BN $(G, \mathcal{P}(G, \Theta(G)))$.

Assumption 1. *The KL divergence between the conditional distributions $\mathcal{P}_i(\pi_i(G), \Theta)$ and $\mathcal{Q}_0 \stackrel{\text{def}}{=} \mathcal{P}_i(\emptyset, \Theta)$ over the variable X_i is bounded by a constant, which for exponential family distribution translates to:*

$$\sup_{\Theta \in \boldsymbol{\varphi}(\mathcal{G})} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)] \leq \Delta_{\max}, \forall i \in [m],$$

where $\Delta_{\max} > 0$ is a constant.

We show that under certain parameterizations, Assumption 1 holds for many commonly used BNs under very mild restrictions on the parameter space $\boldsymbol{\varphi}(\mathcal{G})$.

Theorem 2. *Let $\mathbf{X} = \{X_1, \dots, X_m\}$ be a set of random variables with $X_i \in \mathcal{X}$. Let G be a DAG structure over \mathbf{X} drawn from some family \mathcal{G} of DAG structures. Further assume that we are given a data set S of n i.i.d samples drawn from a BN $(G, \mathcal{P}(G, \Theta(G)))$, where the parameter map Θ , is drawn from some family $\boldsymbol{\varphi}(\mathcal{G})$. Assuming that the condition in Assumption 1 holds for the conditional distributions \mathcal{P}_i , then If $n \leq \mathcal{L}(\mathcal{G})$, then any decoder ζ fails to recover the true DAG structure G with probability $p_{\text{err}} \geq 1/2$, where p_{err} is defined according to (1). The necessary number of samples $\mathcal{L}(\mathcal{G})$, for various classes of BNs, is given as follows:*

$$\mathcal{L}(\tilde{\mathcal{G}}_m) = \frac{\log 2}{\Delta_{\max}} \left(\frac{m-3}{2} - \frac{1}{m} \right),$$

$$\begin{aligned}\mathcal{L}(\tilde{\mathcal{G}}_{m,k}) &= \frac{k \{ \log(m-2)! - \log k! - (m-k-2) \log k \} + \{ (k(k-3)/2) - 1 \} \log 2}{2m\Delta_{\max}}, \\ \mathcal{L}(\mathcal{G}_m^l) &= \frac{\log 2}{2(m-m_l)\Delta_{\max}} \left(\sum_{i=1}^{l-1} m_{i+1}m_i - 2 \right), \\ \mathcal{L}(\mathcal{G}_{m,k}^l) &= \frac{k \sum_{i=1}^{l-1} m_i \log \left(\frac{m_{i+1}}{k} \right) - 2 \log 2}{2(m-m_l)\Delta_{\max}}.\end{aligned}$$

Since the difficulty of learning a class of BNs is determined by the difficulty of learning the most difficult subset within that class, we immediately get the following corollary on the fundamental limits of learning non-sparse and sparse BNs.

Corollary 2 (Fundamental limits of BN structure learning). *The necessary number of samples required to learn BNs on m variables and sparse BNs on m variables and maximum number of parents k is as follows:*

$$\begin{aligned}\mathcal{L}(\mathcal{G}_m) &= \frac{\log 2}{\Delta_{\max}} \left(\frac{m-3}{2} - \frac{1}{m} \right), \\ \mathcal{L}(\mathcal{G}_{m,k}) &= \frac{1}{2\Delta_{\max}} \left(k \log(m-2) + \frac{k(k-3) \log 2}{2m} - R(m,k) \right),\end{aligned}$$

where $R(m,k) = \binom{k}{m} \{ (m-2) + 2 \log(m-2) + \log k! + (m-k-2) \log k \} + (\log 2)/m$.

In the above corollary, the lower bounds for non-sparse and sparse BNs come from the results for the restricted ensembles $\tilde{\mathcal{G}}_m$ and $\tilde{\mathcal{G}}_{m,k}$ respectively.

Remark 4. For sparse BNs, if k is constant, then the reminder term in the lower bound is $R(m,k) \leq 1$ for sufficiently large m and the sample complexity grows as $\Omega(\log m)$. In this regime, our lower bound for learning sparse Bayesian networks matches the analogous lower bound of $\mathcal{O}(\log m)$ for learning degree-bounded Ising models as obtained by Santhanam and Wainwright [6]. In general, however, $R(m,k) = \mathcal{O}(k \log k)$; therefore the number of samples necessary to learn sparse BNs grows as $\Omega(k \log m + k^2/m)$. This is slightly weaker than the lower bound of $\Omega(k^2 \log m)$ for learning degree-bounded Ising models [6].

Remark 5. Note that since the number of samples required to learn general essential DAGs and layered networks is asymptotically similar, i.e. $\mathcal{L}(\tilde{\mathcal{G}}_m) = \mathcal{L}(\mathcal{G}_m^l) = \Omega(m)$. This, combined with the fact that our lower bounds on the number of DAG structures in the ensemble $\tilde{\mathcal{G}}_m$ was tight, leads us to the conclusion that the ordering of variables does not add much to the difficulty of learning BNs.

Theorem 2 provides a simple recipe for obtaining necessary conditions on the sample complexity of learning any exponential family BN, as we demonstrate in the next section.

6 Implications for Commonly Used Bayesian Networks

In this section, we instantiate Theorem 2 for specific conditional distributions, to derive fundamental limits of learning various widely used BNs. This also allows us to highlight the role of parameters of the conditional distributions in the sample complexity of learning the DAG structure of BNs. Table 1 summarizes our results for various commonly used BNs.

	Non-Sparse	Sparse
CPT	$\frac{m}{\log(1/\theta_{\min})}$	$\frac{k \log m + k^2/m}{\log(1/\theta_{\min})}$
Gaussian	$\frac{\sigma_{\min}^2 m}{\sigma_{\min}^2 + 2\mu_{\max}^2(w_{\max}^2 + 1)}$	$\frac{\sigma_{\min}^2 (k \log m + k^2/m)}{\sigma_{\min}^2 + 2(w_{\max}^2 + 1)w_{\max}^2}$
Noisy-OR	$\frac{m}{ \log(\hat{\theta}/(1-\hat{\theta})) }$	$\frac{k \log m + k^2/m}{ \log(\hat{\theta}/(1-\hat{\theta})) }$
Logistic	$\frac{m}{w_{\max}^1}$	$\frac{k \log m + k^2/m}{w_{\max}^1}$

Table 1: Fundamental limits of learning the structure of various types of BNs from Corollary 2, where the entries of the tables are lower bounds i.e. $\Omega(\cdot)$. For CPT BNs, θ_{\min} is the minimum entry in the conditional probability tables. For Gaussian BNs, w_{\max}, μ_{\max} and σ_{\min} are the maximum ℓ_2 norm of the weight vectors, maximum absolute mean and minimum conditional variance respectively. For noisy-OR networks, $\hat{\theta} \in (0, 1)$ is the failure probability. Lastly, for logistic regression, w_{\max}^1 is the maximum ℓ_1 norm of the weight vectors.

Conditional Probability Table BNs. Conditional Probability Table (CPT) based BNs are perhaps the most widely used BNs, where the conditional distribution of a node given its parents is described by probability tables. As is typically the case, we assume that the support of $X_i \in \mathcal{X} = [v]$ for all i . The conditional distribution of X_i is given by the following categorical distribution:

$$\mathcal{P}_i(x_i; \pi_i(G), \Theta) = \prod_{j=1}^v (\theta_{ij}(\mathbf{x}))^{1_{[X_i=j]}},$$

where $\Theta(G) \in \varphi(\mathcal{G})$ are the set of conditional probability tables for the variables $\{X_1, \dots, X_m\}$ compatible with the DAG structure G . Let us denote $\Theta^G \stackrel{\text{def}}{=} \Theta(G)$. The conditional probability table for the i -th random variable $\Theta_i^G : [v]^{|\pi_i|} \rightarrow \Delta_v$, maps all possible assignments to the parent set \mathbf{X}_{π_i} to the $(v-1)$ -dimensional probability simplex Δ_v , and $\theta_{ij}^G(\cdot)$ represents the j -th entry of the v -dimensional vector $\Theta_i^G(\cdot)$. The following lemma gives the upper bound on the mutual information $I(\mathbf{S}; G|\Theta)$ for CPT BNs.

Lemma 7 (Mutual Information bound for CPT networks). *For CPT BNs we have*

$$\begin{aligned} \Delta_{\max} &\leq 4 \log(1/\theta_{\min}), \\ \sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) &\leq 4nm \log(1/\theta_{\min}), \end{aligned}$$

where $\theta_{\min} > 0$ is minimum probability value in a probability table across all node and parent set assignments i.e.,

$$\theta_{\min} \stackrel{\text{def}}{=} \inf_{\Theta \in \varphi(\mathcal{G})} \min_{G \in \mathcal{G}} \min_{i=1}^m \min_{\mathbf{x} \in \mathcal{X}^{|\pi_i(G)|}} \min_{j=1}^v \theta_{ij}^G(\mathbf{x}).$$

Remark 6. The necessary number of samples required to learn dense and sparse CPT BNs is $\Omega\left(\frac{m}{\log(1/\theta_{\min})}\right)$ and $\Omega\left(\frac{k \log m + k^2/m}{\log(1/\theta_{\min})}\right)$, respectively. In the regime that $\theta_{\min} \geq \exp(-1/m)$, the sample complexity for learning dense and sparse BNs is $\Omega(m^2)$ and $\Omega(km \log m + k^2)$, respectively.

Remark 7. The sample complexity of SparsityBoost algorithm by Brenner and Sontag [3] for recovering the structure of binary-valued, sparse, CPT BNs grows as $\mathcal{O}(\max((\log m)\mu, m^2\hat{\mu}_P^2))$, where μ is defined in [3] as “the maximum inverse probability of an assignment to a separating set over all pairs of nodes”, and is $1/\theta_{\min}$ for the ensembles we consider. The parameter $\hat{\mu}_P^2$ is also defined as the maximum inverse probability of an assignment to a separating set but relates to the true graph, G , that generated the data and can be $\ll 1/\theta_{\min}$. If SparsityBoost operates in the regime where the second term inside the max function dominates, which the authors believe to be the case, then that leads to a sufficient condition of $\mathcal{O}((1/\theta_{\min})m^2)$, which is quite far from the information-theoretic limit.

Gaussian BNs. In this case, we assume that the support $X_i \in \mathcal{X} = \mathbb{R}$ for all i , the parameters of the i -th node $\Theta_i(G) = (\mathbf{w}_i^G, \mu, \sigma^2)$, and the conditional distributions are described by the following linear Gaussian model:

$$\mathcal{P}_i(\pi_i(G), \Theta) = \mathcal{N}(\mu_i, \sigma^2/2), \quad (14)$$

$$\mu_i = \begin{cases} (\mathbf{w}_i^G)^T \mathbf{X}_{\pi_i} & \pi_i(G) \neq \emptyset, \\ \mu & \text{otherwise} \end{cases}, \quad (15)$$

where $t_G : [m] \rightarrow [m]$ is a function that maps a node to its “topological order” as defined by the graph G . We assume that

$$\begin{aligned} \mathbf{w}_i^G &\in \mathbb{B}_i^G \stackrel{\text{def}}{=} \{\mathbf{w} \in \mathbb{R}^{|\pi_i(G)|} \mid \|\mathbf{w}\|_2 \leq 1/\sqrt{2(t_G(i)-1)}\}, \\ \mu &\in [\mu_a, \mu_b], \sigma \in [\sigma_{\min}, \sigma_{\max}], \text{ and} \\ \Theta(G) &\in (\times_{i=1}^m \mathbb{B}_i^G) \times [\mu_a, \mu_b] \times [\sigma_{\min}, \sigma_{\max}], \end{aligned}$$

where $-\infty < \mu_a \leq \mu_b < \infty$, $0 < \sigma_{\min} \leq \sigma_{\max} < \infty$. Accordingly, we have that $\mathcal{Q}_0 = \mathcal{P}_i(\emptyset, \Theta) = \mathcal{N}(\mu, \sigma^2/2)$, and $\mu_i \leq \mu \|\mathbf{w}_i^G\|_2$. Once again, we first bound the mutual information $I(\mathbf{S}; G|\Theta)$ in the following lemma which we then plugin in Theorem 2 to obtain the necessary conditions for learning Gaussian BNs.

Lemma 8 (Mutual Information bound for Gaussian networks). *For Gaussian BNs we have:*

$$\begin{aligned} \Delta_{\max} &\leq 1 + \frac{2\mu_{\max}^2(w_{\max}^2 + 1)}{\sigma_{\min}^2}, \\ \sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) &\leq nm \left(1 + \frac{2\mu_{\max}^2(w_{\max}^2 + 1)}{\sigma_{\min}^2} \right), \end{aligned}$$

where $\mu_{\max} \stackrel{\text{def}}{=} \max(|\mu_a|, |\mu_b|)$ and w_{\max} is the maximum ℓ_2 norm of the weight vectors, i.e.

$$w_{\max} \stackrel{\text{def}}{=} \sup_{\Theta \in \varphi(\mathcal{G})} \max_{G \in \mathcal{G}} \max_{i=1}^m \|\mathbf{w}_i^G\|_2.$$

Remark 8. Invoking Theorem 2 for a two layer ordering of nodes, where there are $m-1$ nodes in the top layer and 1 node in the bottom layer, we recover the information-theoretic limits of linear regression. Specifically, we have that the necessary number of samples for linear regression and sparse linear regression scale as $\Omega\left(\frac{\sigma_{\min}^2 m}{\sigma_{\min}^2 + 2\mu_{\max}^2 w_{\max}^2}\right)$ and $\Omega\left(\frac{\sigma_{\min}^2 k \log m}{\sigma_{\min}^2 + 2\mu_{\max}^2 w_{\max}^2}\right)$ respectively.

Remark 9. The number of samples necessary for learning the structure of degree bounded Gaussian MRFs scales as $\Omega(k \log(m/k)/\log(1+k\lambda))$ [7], where λ is the minimum correlation between pairs of nodes that are connected by an edge. The corresponding result for sparse BNs is $\Omega\left(\frac{\sigma_{\min}^2(k \log m + k^2/m)}{\sigma_{\min}^2 + 2\mu_{\max}^2 w_{\max}^2}\right)$, which is slightly stronger than the corresponding lower bound for learning Gaussian MRFs, with respect to sparsity index k .

Noisy-OR BNs. Noisy-OR BNs are another widely used class of BNs — a popular example being the two-layer QMR-DT network [8]. They are usually parameterized by failure probabilities θ_{ij} , which in the context of the QMR-DT network of diseases and symptoms can be interpreted as the probability of not observing the i -th symptom given that the j -th disease is present. More formally, we have binary valued random variables i.e. $X_i \in \mathcal{X} = \{0, 1\}$ for all i , the i -th conditional distribution is given by the Bernoulli distribution (\mathcal{B}) with parameter $\Theta_i = \theta \in (0, 1)$:

$$\mathcal{P}_i(\pi_i(G), \Theta_i) = \mathcal{B}(1 - \theta_i) \quad \mathcal{P}_i(\emptyset, \Theta_i) = \mathcal{B}(\theta), \quad (16)$$

where $\theta_i = \theta \left(\prod_{j \in \pi_i} \theta^{X_j} \right)^{1/|\pi_i|}$. The following lemma bounds the mutual information for noisy-OR networks.

Lemma 9 (Mutual Information bound for Noisy-OR). *For Noisy-OR BNs we have:*

$$\Delta_{\max} \leq 2 \left| \log(\hat{\theta}/(1-\hat{\theta})) \right| \\ \sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) \leq 2nm \left| \log(\hat{\theta}/(1-\hat{\theta})) \right|,$$

where $\hat{\theta} \stackrel{\text{def}}{=} \operatorname{argmax}_{\theta \in \varphi(\mathcal{G})} |\log(\theta/(1-\theta))|$.

Remark 10. From the above lemma we notice that recovering the structure of noisy-OR networks becomes more difficult as the failure probability θ moves farther away from $1/2$. That is because as $\theta \rightarrow 1$, the noisy-OR network becomes more “noisy”. While, as $\theta \rightarrow 0$, the top level nodes (nodes with no parents) take the value 1 with low probability, in which case the child nodes take values 0 with high probability.

Logistic Regression BNs. For logistic regression BNs, the nodes are assumed to be binary valued i.e. $X_i \in \mathcal{X} = \{0, 1\}$ for all i . Each node in the network can be thought of as being classified as “0” or “1” depending on some linear combination of the values of its parents. The parameters for the i -th conditional distribution are $\Theta_i = \mathbf{w}_i$, where the vectors \mathbf{w}_i are assumed to have bounded ℓ_1 norm, i.e. $\mathbf{w}_i \in \mathbb{R}^{|\pi_i|} \wedge \|\mathbf{w}_i\|_1 \leq w_{\max}^1$, for some constant w_{\max}^1 . The conditional distribution of the nodes are given as:

$$\mathcal{P}_i(\pi_i(G), \Theta_i) = \mathcal{B}(\sigma(\langle \mathbf{X}_{\pi_i}, \mathbf{w}_i \rangle)), \quad \mathcal{P}_i(\emptyset, \Theta_i) = \mathcal{B}(1/2), \quad (17)$$

where \mathcal{B} is the Bernoulli distribution and $\sigma(x) = 1/(1 + \exp(-x))$ is the sigmoid function. The following lemma upper bounds the mutual information for logistic regression BNs.

Lemma 10 (Mutual Information bound for Logistic regression networks). *For Logistic regression BNs we have:*

$$\Delta_{\max} \leq w_{\max}^1/2, \quad \sup_{\Theta \in \varphi(\mathcal{G})} I(\mathbf{S}; G|\Theta) \leq (nmw_{\max}^1)/2,$$

where $w_{\max}^1 \stackrel{\text{def}}{=} \sup_{\Theta \in \varphi(\mathcal{G})} \max_{i \in [m]} \|\Theta_i(G)\|_1$.

Remark 11. Once again, we can instantiate Theorem 2 for the two-layer case and obtain necessary number of samples for support recovery in logistic regression. We have that the number of samples needed for support recovery in logistic regression scales as $\Omega(k \log(m)/w_{\max}^1)$. In the regime that $w_{\max}^1 \leq 1/k$, the necessary number of samples scales as $\Omega(k^2 \log m)$. Ravikumar et al. [4] studied support recovery in logistic regression in the context of learning sparse Ising models. The upper bound of $\mathcal{O}(k^2 \log m)$ in Proposition 1 in [4], is thus information-theoretically optimal.

7 Concluding Remarks

An important direction for future work is to study the information-theoretic limits of both structure and parameter recovery of BNs. However, the analysis for that situation is complicated by the fact that one has to come up with an appropriate joint distribution on the structures and parameters of the ensembles. While it is possible to do so for BNs with specific conditional distributions, we anticipate that coming up with general results for BNs would be hard, if at all possible. Also of complimentary interest is the problem of obtaining sharp thresholds for structure learning of Bayesian networks. However, such analysis might also need to be done on a case-by-case basis for specific BNs.

References

- [1] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [2] Cassio P De Campos and Qiang Ji. Efficient structure learning of bayesian networks using constraints. *The Journal of Machine Learning Research*, 12:663–689, 2011.
- [3] Eliot Brenner and David Sontag. SparsityBoost: A New Scoring Function for Learning Bayesian Network Structure. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 112–121, Corvallis, Oregon, 2013. AUAI Press.
- [4] Pradeep Ravikumar, Martin Wainwright, and John Lafferty. High-dimensional Ising model selection using L1-regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- [5] Steen A Andersson, David Madigan, and Michael D Perlman. On the markov equivalence of chain graphs, undirected graphs, and acyclic digraphs. *Scandinavian Journal of Statistics*, 24(1):81–102, 1997.
- [6] Narayana P Santhanam and Martin J Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *Information Theory, IEEE Transactions on*, 58(7):4117–4134, 2012.
- [7] Wei Wang, Martin J Wainwright, and Kannan Ramchandran. Information-theoretic bounds on model selection for gaussian markov random fields. In *ISIT*, pages 1373–1377. Citeseer, 2010.
- [8] Michael A Shwe, B Middleton, DE Heckerman, M Henrion, EJ Horvitz, HP Lehmann, and GF Cooper. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base. *Methods of information in Medicine*, 30(4):241–255, 1991.
- [9] Bertran Steinsky. Enumeration of labelled chain graphs and labelled essential directed acyclic graphs. *Discrete Mathematics*, 270(1-3):267–278, 2003.
- [10] Klaus-U Höffgen. Learning and robust learning of product distributions. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 77–83. ACM, 1993.

- [11] Nir Friedman and Zohar Yakhini. On the sample complexity of learning bayesian networks. In *Proceedings of the Twelfth international conference on Uncertainty in Artificial Intelligence*, pages 274–282. Morgan Kaufmann Publishers Inc., 1996.
- [12] P. Abbeel, D. Koller, and A. Ng. Learning factor graphs in polynomial time and sample complexity. *UAI*, 2005.
- [13] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*, volume 81. MIT press, 2000.
- [14] J. Cheng, R. Greiner, J. Kelly, D. Bell, and W. Liu. Learning Bayesian networks from data: An information-theory based approach. *Artificial Intelligence Journal*, 2002.
- [15] D. Chickering and C. Meek. Finding optimal Bayesian networks. *UAI*, 2002.
- [16] O. Zuk, S. Margel, and E. Domany. On the number of samples needed to learn the correct structure of a Bayesian network. *UAI*, 2006.
- [17] Animashree Anandkumar, Vincent YF Tan, Furong Huang, and Alan S Willsky. High-dimensional gaussian graphical model selection: Walk summability and local separation criterion. *The Journal of Machine Learning Research*, 13(1):2293–2337, 2012.
- [18] Animashree Anandkumar, Vincent YF Tan, Furong Huang, Alan S Willsky, et al. High-dimensional structure estimation in ising models: Local separation criterion. *The Annals of Statistics*, 40(3):1346–1375, 2012.
- [19] R W Robinson. Counting unlabeled acyclic digraphs. *Combinatorial Mathematics V*, 622:28–43, 1977.

Appendix

Proofs for Main Results

Proof of Theorem 1 (Fano's inequality extension). Let,

$$E \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } X \neq \hat{X} \\ 0 & \text{otherwise} \end{cases},$$

and $p_{\text{err}} \stackrel{\text{def}}{=} \Pr\{X \neq \hat{X}\}$. Then using the chain rule for entropy we can expand the conditional entropy $H(E, X|\hat{X}, W)$ as follows:

$$H(E, X|\hat{X}, W) = H(E|X, \hat{X}, W) + H(X|\hat{X}, W) \quad (18)$$

$$= H(X|E, \hat{X}, W) + H(E|\hat{X}, W) \quad (19)$$

Next, we bound each of the terms in (18) and (19). $H(E|X, \hat{X}, W) = 0$ because E is a deterministic function of X and \hat{X} . Moreover, since conditioning reduces entropy, we have that $H(E|\hat{X}, W) \leq H(E) = H(p_{\text{err}})$. Using the same arguments we have the following upper bound on $H(X|E, \hat{X}, W)$:

$$\begin{aligned} H(X|E, \hat{X}, W) &= p_{\text{err}}H(X|E=1, \hat{X}, W) + (1-p_{\text{err}})H(X|E=0, \hat{X}, W) \\ &\leq p_{\text{err}}H(X|W) \end{aligned} \quad (20)$$

Next, we show that $I(\hat{X}; X|W) \leq I(\hat{X}; Y|W)$, which can be thought of as the conditional data processing inequality. Using the chain rule of mutual information we have that

$$\begin{aligned} I(Y, \hat{X}; X|W) &= I(\hat{X}; X|Y, W) + I(Y; X|W) \\ &= I(Y; X|\hat{X}, W) + I(\hat{X}; X|W). \end{aligned}$$

Since, conditioned on Y , X and \hat{X} are independent. We have $I(\hat{X}; X|Y, W) = 0$.

$$\begin{aligned} \implies I(Y; X|W) &= I(Y; X|\hat{X}, W) + I(\hat{X}; X|W) \\ \implies I(Y; X|W) &\geq I(\hat{X}; X|W). \end{aligned}$$

Therefore, we can bound $H(X|\hat{X}, W)$ as follows:

$$H(X|\hat{X}, W) = H(X|W) - I(\hat{X}; X|W) \geq H(X|W) - I(Y; X|W). \quad (21)$$

Combining (18),(19),(20) and (21), we get:

$$\begin{aligned} H(X|W) - I(Y; X|W) &\leq H(p_{\text{err}}) + p_{\text{err}}H(X|W) \\ \implies H(X|W) - I(Y; X|W) &\leq \log 2 + p_{\text{err}}H(X|W) \\ \implies p_{\text{err}} &\geq 1 - \frac{I(Y; X|W) + \log 2}{H(X|W)} \end{aligned} \quad (22)$$

Now if X and W are independent then $H(X|W) = H(X)$. Denoting the joint distribution over X, Y, W by $\mathcal{P}_{X,Y,W}$, the conditional distribution of X, Y given W by $\mathcal{P}_{X,Y|W}$ and so on; the final claim follows from bounding the term $I(Y; X|W)$ as follows:

$$\begin{aligned} I(Y; X|W) &= \mathbb{E}_{\mathcal{P}_{X,Y,W}} \left[\log \frac{\mathcal{P}_{X,Y|W}}{\mathcal{P}_{X|W}\mathcal{P}_{Y|W}} \right] = \mathbb{E}_{\mathcal{P}_W} \left[\mathbb{E}_{\mathcal{P}_{X,Y|W=w}} \left[\log \frac{\mathcal{P}_{X,Y|W=w}}{\mathcal{P}_{X|W=w}\mathcal{P}_{Y|W=w}} \right] \right] \\ &\leq \sup_{w \in \mathcal{W}} I(X; Y|W = w). \end{aligned}$$

□

Proof of Lemma 1. First, we briefly review Steinsky's method for counting essential DAGs, which in turn is based upon Robinson's [19] method for counting labeled DAGs. The main idea behind Steinsky's method is to split the set of essential DAGs into overlapping subsets with different terminal vertices — vertices with out-degree 0. Let $A_i \subset \tilde{\mathcal{G}}_m$ be the set of essential DAGs where the i -th node is a terminal node. Using the inclusion-exclusion principle, the number of essential DAGs is given as follows:

$$\begin{aligned} c_m &\stackrel{\text{def}}{=} |\tilde{\mathcal{G}}_m| = |A_1 \cup \dots \cup A_m| \\ &= \sum_{s=1}^m (-1)^{(s+1)} \sum_{1 \leq i_1 \leq \dots \leq i_s \leq m} |A_{i_1} \cap \dots \cap A_{i_s}|. \end{aligned} \quad (23)$$

Now, consider the term $|A_1 \cap \dots \cap A_{m-1}|$, i.e., number of essential DAGs where nodes $[m-1]$ are terminal nodes. The number of ways of adding the m -th vertex as a terminal vertex to an arbitrary essential DAG on the nodes $[m-1]$ is: $2^{m-1} - (m-1)$. The term $m-1$ needs to be subtracted to account for edges that are not protected. Therefore, c_m is given by the following recurrence relation:

$$c_m = \sum_{s=1}^m (-1)^{s+1} \binom{m}{s} (2^{m-s} - (m-s))^s c_{m-s}, \quad (24)$$

where $c_0 = 1$. Using Bonferroni's inequalities we can upper bound c_m as follows:

$$\begin{aligned} c_m &\leq m(2^{m-1} - (m-1))c_{m-1} \leq m2^{m-1}c_{m-1} \\ &\leq m! 2^{m(m-1)/2}. \end{aligned} \quad (25)$$

Using Bonferroni's inequalities to lower bound c_m produces recurrence relations that have no closed form solution. Therefore, we lower bound c_m in (23) as follows:

$$\begin{aligned} c_m &= |A_1 \cup \dots \cup A_m| \geq \max_i |A_i| \\ &= (2^{m-1} - (m-1))c_{m-1} \geq 2^{m-2}c_{m-1} \\ &\geq 2^{(m(m-3)/2)+1}. \end{aligned} \quad (26)$$

□

Proof of Lemma 2. In this case, A_i is the set of essential DAGs where the i -th node is a terminal node and all nodes have at most k parents. Once again, using the inclusion-exclusion principle, the number of essential DAGs with at most k parents is given as follows:

$$c_{m,k} \stackrel{\text{def}}{=} |\mathcal{G}_{m,k}| = \sum_{s=1}^m (-1)^{(s+1)} \sum_{1 \leq i_1 \leq \dots \leq i_s \leq m} |A_{i_1} \cap \dots \cap A_{i_s}| \quad (27)$$

Now, consider the term $|A_1 \cap \dots \cap A_s|$, i.e. number of essential DAGs where nodes $\{1, \dots, s\}$ are terminal nodes. Let G_0 be an arbitrary essential DAG over nodes $\{s+1, \dots, m\}$, where each node has at most k parents. Let $u \in \{1, \dots, s\}$ and $v \in \{s+1, \dots, m\}$ be arbitrary nodes. Let G_1 be the new graph, formed by connecting u to G_0 . The edge $v \rightarrow u$ is not protected, or covered, in G_1 if $\pi_v(G_0) = \pi_u(G_1) \setminus \{v\}$. For each node $v \in \{s+1, \dots, m\}$, there is exactly one configuration in which the edge $v \rightarrow u$ is covered, i.e. when we set the parents of u to be $\pi_v(G_0) \cup \{v\}$; unless $|\pi_v(G_0)| = k$, in which case $v \rightarrow u$ is always protected. Let $\kappa(G_0)$ be the

number of nodes in G_0 that have less than k parents. Then the number of ways of adding a terminal vertex u to G_0 is: $\sum_{i=0}^k \binom{m-s}{i} - \kappa(G_0)$ when $(m-s) > k$, and $2^k - k$ when $(m-s) \leq k$. We can simply bound $\kappa(G_0)$ by: $0 \leq \kappa(G_0) \leq m-s$. This gives a lower bound on the number of ways to add a terminal vertex to G_0 as: $\sum_{i=0}^k \left\{ \binom{m-s}{i} - \frac{(m-s)}{k+1} \right\} \geq \sum_{i=0}^k \binom{m-s-1}{i}$. Using the fact that $\max_{i=1}^m |A_i| \leq c_{m,k} \leq \sum_{i=1}^m |A_i|$, we get the following recurrence relation for upper and lower bounds on the number of essential DAGs with at most k parents:

$$c_{m,k} \leq m \left(\sum_{i=0}^k \binom{m-1}{i} \right) c_{m-1,k} \quad (28)$$

$$c_{m,k} \geq \left(\sum_{i=0}^k \binom{m-2}{i} \right) c_{m-1,k}, \quad (29)$$

where from Lemma 1 we have that $2^{(k(k-3)/2)+1} \leq c_{k,k} \leq k! 2^{k(k-1)/2}$. Thus, we can upper bound $c_{m,k}$ as follows:

$$c_{m,k} \leq m! 2^{k(k-1)/2} \prod_{j=k+1}^{m-1} \left(\sum_{i=0}^k \binom{j}{i} \right) \quad (30)$$

Similarly, we can lower bound $c_{m,k}$ as follows:

$$c_{m,k} \geq 2^{(k(k-3)/2)+1} \prod_{j=k+1}^{m-1} \left(\sum_{i=0}^k \binom{j-1}{i} \right) \quad (31)$$

Finally, using (31), we lower bound $\log c_{m,k}$ as follows:

$$\log c_{m,k} \geq ((k(k-3)/2) + 1) \log 2 + \sum_{j=k+1}^{m-1} \log \left(\sum_{i=0}^k \binom{j-1}{i} \right)$$

The second term in the above equation is lower bounded as follows:

$$\begin{aligned} \sum_{j=k+1}^{m-1} \log \left(\sum_{i=0}^k \binom{j-1}{i} \right) &\geq \sum_{j=k}^{m-2} \log \left(\sum_{i=0}^k \binom{j}{i} \right) \geq \sum_{j=k}^{m-2} \log \binom{j}{k} \geq \sum_{j=k}^{m-2} \log \left(\frac{j}{k} \right)^k \\ &\geq k \{ \log(m-2)! - \log k! - (m-k-2) \log k \}. \end{aligned}$$

□

Proof Lemma 3. For layered non-sparse BNs, the number of possible choices for parents of a node in layer i is $2^{m_{i+1}}$. Therefore, the total number of non-sparse Bayesian networks is given as $\prod_{i=1}^{l-1} (2^{m_{i+1}})^{m_i}$. Similarly, for the sparse case, the number of possible choices for parents of a node in layer i is $\sum_{j=0}^k \binom{m_{i+1}}{j}$. Therefore, the total number of sparse Bayesian networks is $\prod_{i=1}^{l-1} \left[\sum_{j=0}^k \binom{m_{i+1}}{j} \right]^{m_i}$. □

Proof of Lemma 4. Let $c \stackrel{\text{def}}{=} |\mathcal{G}|$, for some ensemble of DAGs \mathcal{G} . Denoting the conditional distribution of the data given a specific instance of the parameters Θ by $\mathcal{P}_{S|\Theta}$, we have:

$$\sup_{\Theta \in \varphi(\mathcal{G})} I(S; G|\Theta) = \sup_{\Theta \in \varphi(\mathcal{G})} \frac{1}{c} \sum_{G \in \mathcal{G}} \mathbb{KL}(\mathcal{P}_{S|G,\Theta} \| \mathcal{P}_{S|\Theta}), \quad (32)$$

where in $\mathbb{KL}(\mathcal{P}_{S|G,\Theta} \parallel \mathcal{P}_{S|\Theta})$, Θ and G are specific instances and not random variables. For any distribution \mathcal{Q} over S , we can rewrite $\mathbb{KL}(\mathcal{P}_{S|G,\Theta} \parallel \mathcal{P}_{S|\Theta})$ as follows:

$$\begin{aligned}\mathbb{KL}(\mathcal{P}_{S|G,\Theta} \parallel \mathcal{P}_{S|\Theta}) &= \mathbb{E}_S \left[\log \frac{\mathcal{P}_{S|G,\Theta}}{\mathcal{Q}} \frac{\mathcal{Q}}{\mathcal{P}_{S|\Theta}} \right] \\ &= \mathbb{KL}(\mathcal{P}_{S|G,\Theta} \parallel \mathcal{Q}) - \mathbb{E}_S \left[\log \frac{\mathcal{P}_{S|\Theta}}{\mathcal{Q}} \right],\end{aligned}\tag{33}$$

where the expectation $\mathbb{E}_S[\cdot]$ is with respect to the distribution $\mathcal{P}_{S|G,\Theta}$. Now, $\mathbb{E}_S \left[\log \frac{\mathcal{P}_{S|\Theta}}{\mathcal{Q}} \right]$ can be written as follows:

$$\begin{aligned}\sum_{G \in \mathcal{G}} \mathbb{E}_S \left[\log \frac{\mathcal{P}_{S|\Theta}}{\mathcal{Q}} \right] &= \sum_{G \in \mathcal{G}} \sum_S \Pr\{S|G, \Theta\} \log \frac{\mathcal{P}_{S|\Theta}}{\mathcal{Q}} \\ &= c \sum_S \sum_{G \in \mathcal{G}} \Pr\{G\} \Pr\{S|G, \Theta\} \log \frac{\mathcal{P}_{S|\Theta}}{\mathcal{Q}} \\ &= c \mathbb{KL}(\mathcal{P}_{S|\Theta} \parallel \mathcal{Q}),\end{aligned}\tag{34}$$

where, once again, we emphasize that in $\mathbb{KL}(\mathcal{P}_{S|\Theta} \parallel \mathcal{Q})$, Θ is a particular instance of the parameters and not a random variable. Combining (32), (33) and (34), and using the fact that $\mathbb{KL}(\mathcal{P}_{S|\Theta} \parallel \mathcal{Q}) > 0$, we get

$$\sup_{\Theta \in \varphi(\mathcal{G})} I(S; G|\Theta) \leq \sup_{\Theta \in \varphi(\mathcal{G})} \frac{1}{c} \sum_{G \in \mathcal{G}} \mathbb{KL}(\mathcal{P}_{S|G,\Theta} \parallel \mathcal{Q})\tag{35}$$

□

Proof of Lemma 5 (KL bound for exponential family).

$$\mathbb{KL}(\mathcal{P}_1 \parallel \mathcal{P}_2) = \boldsymbol{\eta}_1^T \mathbb{E}_X [\mathbf{T}(x)|\boldsymbol{\eta}_1] - \psi(\boldsymbol{\eta}_1) - \boldsymbol{\eta}_2^T \mathbb{E}_X [\mathbf{T}(x)|\boldsymbol{\eta}_1] + \psi(\boldsymbol{\eta}_2),\tag{36}$$

where for computing the expected sufficient statistic, $\mathbb{E}_X [\mathbf{T}(x)|\boldsymbol{\eta}_1]$, we take the expectation with respect to the distribution parameterized by $\boldsymbol{\eta}_1$. Now, from the mean value theorem we have that

$$\begin{aligned}\psi(\boldsymbol{\eta}_2) - \psi(\boldsymbol{\eta}_1) &= \nabla \psi(\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1)^T [\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1] \\ &= (\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)^T \mathbb{E}_X [\mathbf{T}(x)|\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1]\end{aligned}$$

for some $\alpha \in [0, 1]$. Then we have that,

$$\begin{aligned}\mathbb{KL}(\mathcal{P}_1 \parallel \mathcal{P}_2) &= \mathcal{T}(\boldsymbol{\eta}_1)^T [\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2] + (\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)^T \mathcal{T}(\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1) \\ &= (\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2)^T \{\mathcal{T}(\boldsymbol{\eta}_1) - \mathcal{T}(\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1)\}\end{aligned}\tag{37}$$

Since the function \mathcal{T} is the gradient of the convex function ψ , it is monotonic. Therefore, the function $\mathcal{T}(\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1)$ takes the maximum value at the end points $\alpha = 0$ or at $\alpha = 1$. Assuming \mathcal{T} is maximized at $\alpha = 0$, the i -th KL divergence term can be upper bound using (37) as:

$$0 \leq \mathbb{KL}(\mathcal{P}_1 \parallel \mathcal{P}_2) \leq (\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2)^T \{\mathcal{T}(\boldsymbol{\eta}_1) - \mathcal{T}(\boldsymbol{\eta}_2)\}$$

On the other hand, assuming \mathcal{T} is maximized at $\alpha = 1$, the i -th KL divergence term can be upper bound using (37) as:

$$0 \leq \mathbb{KL}(\mathcal{P}_1 \| \mathcal{P}_2) \leq 0.$$

Therefore, clearly, $\mathcal{T}(\alpha \boldsymbol{\eta}_2 + (1 - \alpha) \boldsymbol{\eta}_1)$ is maximized at $\alpha = 0$. \square

Proof of Theorem 2. Setting the measure $\mathcal{P}_{\mathcal{G}}$ to be the uniform over \mathcal{G} , and using the Fano's inequality from Theorem 1 and the mutual information bound from Lemma 6, combined with our Assumption 1, we can bound the estimation error as follows:

$$p_{\text{err}} \geq 1 - \frac{nm\Delta_{\max} + \log 2}{\log |\mathcal{G}|}.$$

Then by using the lower bounds on the number of DAG structures in each of the ensembles from Lemmas 1, 2 and 3, and setting p_{err} to $1/2$, we prove our claim. \square

Proofs for Commonly Used Bayesian Networks

Proof of Lemma 7 (Mutual Information bound for CPT networks). For CPT, the mutual information bound is representative of the case when we do not have a closed form solution for the marginal and joint distributions; yet, we can easily bound $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$ through a simple application of the Cauchy-Schwartz inequality, and obtain tighter bounds than the naive $\mathcal{O}(mn \log v)$ bound on the mutual information $I(\mathbf{S}; G | \Theta)$. The sufficient statistics and the natural parameter for the categorical distribution is given as follows:

$$\mathbf{T}(x) = (\mathbf{1}[x = j])_{j=1}^v, \quad \boldsymbol{\eta}_i(\mathbf{X}_{\pi_i}, \Theta_i) = (\log \theta_{ij}(\mathbf{X}_{\pi_i}))_{j=1}^v.$$

Therefore, the expected sufficient statistic $\mathcal{T}(\boldsymbol{\eta}_i) = \Theta_i(\mathbf{X}_{\pi_i})$. From that we get the following upper bound

$$\begin{aligned} \sup_{\Theta \in \varphi(G)} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)] &= \sup_{\Theta \in \varphi(G)} \mathbb{E}_{\mathbf{X}_{\pi_i}} [(\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)^T \{\mathcal{T}(\boldsymbol{\eta}_i) - \mathcal{T}(\boldsymbol{\eta}_0)\}] \\ &\leq \sup_{\Theta \in \varphi(G)} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\|\boldsymbol{\eta}_i - \boldsymbol{\eta}_0\|_{\infty} \|\mathcal{T}(\boldsymbol{\eta}_i) - \mathcal{T}(\boldsymbol{\eta}_0)\|_1] \\ &\leq 4 \log(1/\theta_{\min}), \end{aligned}$$

where in the above we used the Cauchy-Schwartz inequality followed by the fact $\|\Theta_i(\mathbf{x})\|_1 = 1, \forall \mathbf{x} \in \mathcal{X}$. \square

Proof of Lemma 8 (Mutual Information bound for Gaussian). This exemplifies the case where we have closed form solutions for the joint and marginal distributions, which in this case is Gaussian, and we can compute the expected value of $\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)$. The sufficient statistics and natural parameter for the i -th conditional distribution are given as follows:

$$\mathbf{T}(X_i) = \frac{X_i}{\sigma/\sqrt{2}}, \quad \boldsymbol{\eta}_i = \frac{\mu_i}{\sigma/\sqrt{2}}.$$

Also note that, $\forall i \in [m]$, the marginal expectation $\mathbb{E}[X_i] \leq \mu \|\mathbf{w}_i\|_2$. Therefore, we have that $\mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)] = \mathbb{E}_{\mathbf{X}_{\pi_i}} [2(\mu_i - \mu)^2 / \sigma^2] = 2(\mathbb{E}_{\mathbf{X}_{\pi_i}} [(\mu_i - \mu)^2] + \text{Var}_{\mathbf{X}_{\pi_i}} [\mu_i]) / \sigma^2 \leq 2(\mu^2 (\|\mathbf{w}_i\|_2 - 1)^2 + \text{Var}_{\mathbf{X}_{\pi_i}} [\mu_i]) / \sigma^2$. Hence, we need to upper bound $\text{Var}_{\mathbf{X}_{\pi_i}} [\mu_i]$ in order to upper bound $\mathbb{E}_{\mathbf{X}_{\pi_i}} [\Delta(\cdot)]$. Let $(i)_G \in [m]$ be the i -th node in the topological order defined by the graph G . We use the shorthand notation

(i), when it is clear from context that the i -th node in the topological ordering is intended. Now, from the properties of the Gaussian distribution we know that if the conditional distributions are all Gaussian, then the joint distribution over any subset of \mathbf{X} is Gaussian as well. Let $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ be the covariance matrix for the joint distribution over \mathbf{X} , and similarly $\mathbf{\Sigma}_{(i)} \in \mathbb{R}^{i \times i}$ denote the covariance matrix for the joint distribution over variables $\{X_{(1)}, \dots, X_{(i)}\}$. Let $\bar{\mathbf{w}}_{(i)} \in \mathbb{R}^{i-1}$ be the weight vector defined as follows:

$$\forall j \in [i-1], (\bar{\mathbf{w}}_{(i)})_j = \begin{cases} 0 & \text{if } j \notin \pi_{(i)}, \\ (w_{(i)})_j & \text{otherwise.} \end{cases}$$

Note that $\|\bar{\mathbf{w}}_{(i)}\|_2 = \|\mathbf{w}_{(i)}\|_2 \leq 1/\sqrt{2(i-1)}$. Then, we have that $\text{Var}[\mu_{(i)}] = \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G}$ and $\text{Var}[X_{(i)}] = \text{Var}[\mu_{(i)}] + \sigma^2/2$. Also, for any $j \in [i-1]$, we have that $\text{Cov}[X_{(i)}X_{(j)}] = \bar{\mathbf{w}}_{(i)G}^T (\mathbf{\Sigma}_{(i-1)G})_{*,j}$, where $(\mathbf{\Sigma}_{(i-1)G})_{*,j}$ is the j -th column of the matrix $\mathbf{\Sigma}_{(i-1)}$. Therefore, the covariance matrix $\mathbf{\Sigma}_{(i)G}$ can be written as follows:

$$\mathbf{\Sigma}_{(i)G} = \begin{bmatrix} \mathbf{\Sigma}_{(i-1)G} & \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G} \\ \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} & \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G} + \sigma^2/2 \end{bmatrix},$$

where $\mathbf{\Sigma}_{(1)} \in \mathbb{R}^{1 \times 1} = [[\sigma^2/2]]$. Since $\mathbf{\Sigma}_{(i)}$ is positive definite, we have that $\lambda_{\max}(\mathbf{\Sigma}_{(i)}) \leq \lambda_{\max}(\mathbf{\Sigma}_{(i-1)}) + \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G} + \sigma^2/2$. Next, we prove, by induction, that $\lambda_{\max}(\mathbf{\Sigma}_{(i)}) \leq i\sigma^2$. First, note that the base case holds, i.e. $\lambda_{\max}(\mathbf{\Sigma}_{(1)}) = \sigma^2/2 \leq \sigma^2$. Now assume, that $\lambda_{\max}(\mathbf{\Sigma}_{(i-1)}) \leq (i-1)\sigma^2$. Then, we have:

$$\begin{aligned} \lambda_{\max}(\mathbf{\Sigma}_{(i)}) &\leq \lambda_{\max}(\mathbf{\Sigma}_{(i-1)}) + \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G} + \frac{\sigma^2}{2} \\ &\leq (i-1)\sigma^2 + \|\bar{\mathbf{w}}_{(i)}\|_2^2 (i-1)\sigma^2 + \frac{\sigma^2}{2} \\ &\leq (i-1)\sigma^2 + \frac{1}{2(i-1)}(i-1)\sigma^2 + \frac{\sigma^2}{2} \\ &\leq i\sigma^2. \end{aligned}$$

Therefore, we can bound the variance of μ_i as follows:

$$\begin{aligned} \text{Var}[\mu_{(i)}] &= \bar{\mathbf{w}}_{(i)G}^T \mathbf{\Sigma}_{(i-1)G} \bar{\mathbf{w}}_{(i)G} \leq \|\bar{\mathbf{w}}_{(i)}\|_2^2 \lambda_{\max}(\mathbf{\Sigma}_{(i-1)}) \\ &\leq \frac{1}{2(i-1)}(i-1)\sigma^2 = \frac{\sigma^2}{2}. \end{aligned}$$

Thus, we have that $\text{Var}[\mu_i] \leq \sigma^2/2$, $\mathbb{E}_{\mathbf{X}_{\pi_i}}[\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)] \leq 1 + 2\mu^2(\|\mathbf{w}_i\|_2 - 1)^2/\sigma^2$, and $\Delta_{\max} \leq 1 + (2\mu_{\max}^2(w_{\max}^2 + 1))/\sigma_{\min}^2$. □

Proof of Lemma 9 (Mutual Information bound for Noisy-OR). The expected sufficient statistics and natural parameter for the Bernoulli distribution is given as:

$$\mathcal{T}(\boldsymbol{\eta}_i) = 1 - \theta_i, \quad \boldsymbol{\eta}_i = \log \frac{\theta_i}{1 - \theta_i}.$$

Also, $\mathcal{T}(\boldsymbol{\eta}_0) = \theta$ and $\boldsymbol{\eta}_0 = \log((1-\theta)/\theta)$. Using the fact that $\theta^2 \leq \theta_i \leq \theta$, we can bound $\mathbb{E}_{\mathbf{X}_{\pi_i}}[\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0)]$ as follows:

$$\Delta(\boldsymbol{\eta}_i, \boldsymbol{\eta}_0) = \left(\log \frac{\theta_i}{1 - \theta_i} - \log \frac{1 - \theta}{\theta} \right) (1 - \theta_i - \theta)$$

$$\begin{aligned}
&\leq \left| \log \frac{\theta_i \theta}{(1 - \theta_i)(1 - \theta)} \right| \\
&\leq 2 \left| \log \frac{\theta}{1 - \theta} \right|
\end{aligned}$$

Therefore, we have that $\Delta_{\max} \leq 2 \left| \log(\hat{\theta}/(1-\hat{\theta})) \right|$.

□

Proof of Lemma 10 (MI bound for Logistic regression networks). The expected sufficient statistics and the natural parameter are given as follows:

$$\mathcal{T}(\boldsymbol{\eta}_i) = \sigma(\mathbf{w}_i^T \mathbf{X}_{\pi_i}), \quad \boldsymbol{\eta}_i = \log \frac{\sigma(\mathbf{w}_i^T \mathbf{X}_{\pi_i})}{1 - \sigma(\mathbf{w}_i^T \mathbf{X}_{\pi_i})} = \mathbf{w}_i^T \mathbf{X}_{\pi_i}.$$

From the above, we also have that $\mathcal{T}(\boldsymbol{\eta}_0) = 1/2$ and $\boldsymbol{\eta}_0 = 0$. Then, Δ_{\max} is bounded as follows:

$$\begin{aligned}
\Delta_{\max} &= \mathbb{E}_{\mathbf{X}_{\pi_i}} [\mathbf{w}_i^T \mathbf{X}_{\pi_i} (\sigma(\mathbf{w}_i^T \mathbf{X}_{\pi_i}) - 1/2)] \\
&\leq \frac{1}{2} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\mathbf{w}_i^T \mathbf{X}_{\pi_i}] \leq \frac{1}{2} \mathbb{E}_{\mathbf{X}_{\pi_i}} [\|\mathbf{w}_i\|_1 \|\mathbf{X}_{\pi_i}\|_{\infty}] \\
&\leq \frac{\|\mathbf{w}_i\|_1}{2} \leq \frac{w_{\max}^1}{2}.
\end{aligned}$$

□